

## Claims

What is claimed is:

1. Apparatus for recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or more end-user transactions, received at a server in communication with the client workstation, the apparatus comprising:

at least one processor operative to: (i) obtain the one or more remote procedure calls; and (ii) recognize one or more end-user transactions from the one or more remote procedure calls based on training data associated with one or more of the remote procedure calls; and

a memory, coupled to the at least one processor, for storing at least a portion of results associated with the end-user transaction recognition operation.

2. The apparatus of claim 1, wherein the at least one processor is further operative to perform the end-user transaction recognition operation by: (i) segmenting the one or more remote procedure calls into end-user transaction segments; (ii) extracting end-user transaction features for each segment; and (iii) computing the one or more end-user transactions from the extracted end-user transaction features for each segment using a model generated from the training data.

3. The apparatus of claim 2, wherein the model is a Bayes Net model.

4. The apparatus of claim 3, wherein the at least one processor is further operative to perform the end-user transaction computing operation by: (i) assigning a feature value to the corresponding extracted features, the feature value being associated with the Bayes Net model; (ii) finding a posterior probability distribution for each feature value; and (iii) selecting a feature value to represent one or more end-user transactions based on the posterior probability distribution.

5. The apparatus of claim 4, wherein the feature value selection operation comprises selecting the most likely feature value for a given extracted feature to represent one or more end-user transactions.

5 6. The apparatus of claim 4, wherein the selected feature values form a class label representative of one or more end-user transactions.

7. The apparatus of claim 1, wherein at least a portion of the results are compared to a quality metric to determine whether the results are of a sufficient quality.

10 8. Apparatus for generating a model for use in recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or more end-user transactions, received at a server in communication with the client workstation, the apparatus comprising:

15 at least one processor operative to: (i) obtain one or more remote procedure calls labeled with one or more end-user transactions; (ii) compute one or more selected features on which the model is to be trained from the labeled remote procedure calls, the one or more selected features being labeled with the one or more end-user transactions; and (iii) construct the model from the selected features; and

a memory, coupled to the at least one processor, for storing at least a portion of results associated with one or more of the model generating operations.

20 9. The apparatus of claim 8, wherein the one or more labeled remote procedure calls are manually generated.

10. The apparatus of claim 8, wherein the one or more labeled remote procedure calls are automatically generated.

11. The apparatus of claim 8, wherein the at least one processor is further operative to construct the model by: (i) learning dependencies between the computed features to form a model structure; and (ii) learning one or more model parameters given the structure to form the model.

5 12. The apparatus of claim 8, wherein the model is a Bayes Net model.

13. Apparatus for automatically generating training data used to construct a model for use in recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or more end-user transactions, received at a server in communication with the client workstation, the apparatus comprising:

10 at least one processor operative to: (i) cause the client workstation to mark the beginning and end of one or more end-user transactions originating at the client workstation; and (ii) correlating the end-user transaction marks with the one or more remote procedure calls received at the server to generate remote procedure calls labeled with end-user transactions, the labeled remote procedure calls serving as the training data; and

15 a memory, coupled to the at least one processor, for storing at least a portion of results associated with one or more of the training data generating operations.

20 14. The apparatus of claim 13, wherein the at least one processor is further operative to download a marking agent to the client workstation for use in marking the end-user transactions.

15. A method of recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or

more end-user transactions, received at a server in communication with the client workstation, the method comprising the steps of:

obtaining the one or more remote procedure calls; and

5 recognizing one or more end-user transactions from the one or more remote procedure calls based on training data associated with one or more of the remote procedure calls.

10 16. The method of claim 15, wherein the end-user transaction recognition step comprises: (i) segmenting the one or more remote procedure calls into end-user transaction segments; (ii) extracting end-user transaction features for each segment; and (iii) computing the one or more end-user transactions from the extracted end-user transaction features for each segment using a model generated from the training data.

17. The method of claim 16, wherein the model is a Bayes Net model.

15 18. The method of claim 17, wherein the end-user transaction computing step comprises: (i) assigning a feature value to the corresponding extracted features, the feature value being associated with the Bayes Net model; (ii) finding a posterior probability distribution for each feature value; and (iii) selecting a feature value to represent one or more end-user transactions based on the posterior probability distribution.

20 19. The method of claim 18, wherein the feature value selection step comprises selecting the most likely feature value for a given extracted feature to represent one or more end-user transactions.

20. The method of claim 18, wherein the selected feature values form a class label representative of one or more end-user transactions.

21. The method of claim 15, wherein at least a portion of the results are compared to a quality metric to determine whether the results are of a sufficient quality.

22. A method of generating a model for use in recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or more end-user transactions, received at a server in communication with the client workstation, the method comprising the steps of:

obtaining one or more remote procedure calls labeled with one or more end-user transactions;

computing one or more selected features on which the model is to be trained from the labeled remote procedure calls, the one or more selected features being labeled with the one or more end-user transactions; and

constructing the model from the selected features.

23. The method of claim 22, wherein the one or more labeled remote procedure calls are manually generated.

24. The method of claim 22, wherein the one or more labeled remote procedure calls are automatically generated.

25. The method of claim 22, wherein the model constructing step comprises: (i) learning dependencies between the computed features to form a model structure; and (ii) learning one or more model parameters given the structure to form the model.

26. The method of claim 22, wherein the model is a Bayes Net model.

27. A method of automatically generating training data used to construct a model for use in recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or more end-user transactions, received at a server in communication with the client workstation, the method comprising the steps of:

causing the client workstation to mark the beginning and end of one or more end-user transactions originating at the client workstation; and

correlating the end-user transaction marks with the one or more remote procedure calls received at the server to generate remote procedure calls labeled with end-user transactions, the labeled remote procedure calls serving as the training data.

28. The method of claim 27, wherein the causing step comprises downloading a marking agent to the client workstation for use in marking the end-user transactions.

29. An article of manufacture for recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or more end-user transactions, received at a server in communication with the client workstation, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

obtaining the one or more remote procedure calls; and

recognizing one or more end-user transactions from the one or more remote procedure calls based on training data associated with one or more of the remote procedure calls.

30. An article of manufacture for generating a model for use in recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or more end-user transactions, received at a

server in communication with the client workstation, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

obtaining one or more remote procedure calls labeled with one or more end-user transactions;

5           computing one or more selected features on which the model is to be trained from the labeled remote procedure calls, the one or more selected features being labeled with the one or more end-user transactions; and

constructing the model from the selected features.

10           31. An article of manufacture for automatically generating training data used to construct a model for use in recognizing one or more end-user transactions originating at a client workstation from one or more remote procedure calls, corresponding to the one or more end-user transactions, received at a server in communication with the client workstation, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

15           causing the client workstation to mark the beginning and end of one or more end-user transactions originating at the client workstation; and

correlating the end-user transaction marks with the one or more remote procedure calls received at the server to generate remote procedure calls labeled with end-user transactions, the labeled remote procedure calls serving as the training data.

20